

Industrial Strength Pipes Overview

a pipeline toolkit for image processing

Jim Garlick

garlick1@llnl.gov

Center for Application Development and Software Engineering
Lawrence Livermore National Laboratory



Introduction

- Purpose: Provide API for componentizing algorithms and a scheme for executing pipelines on large and small machines.
- Goal: Simple.
- Goal: Lightweight component technology.
- Goal: Abstract data-parallel execution.
- Goal: Enable further middleware research.
- Status: internal “pre-alpha” software running some of SuperMACHO on LLNL Linux clusters.

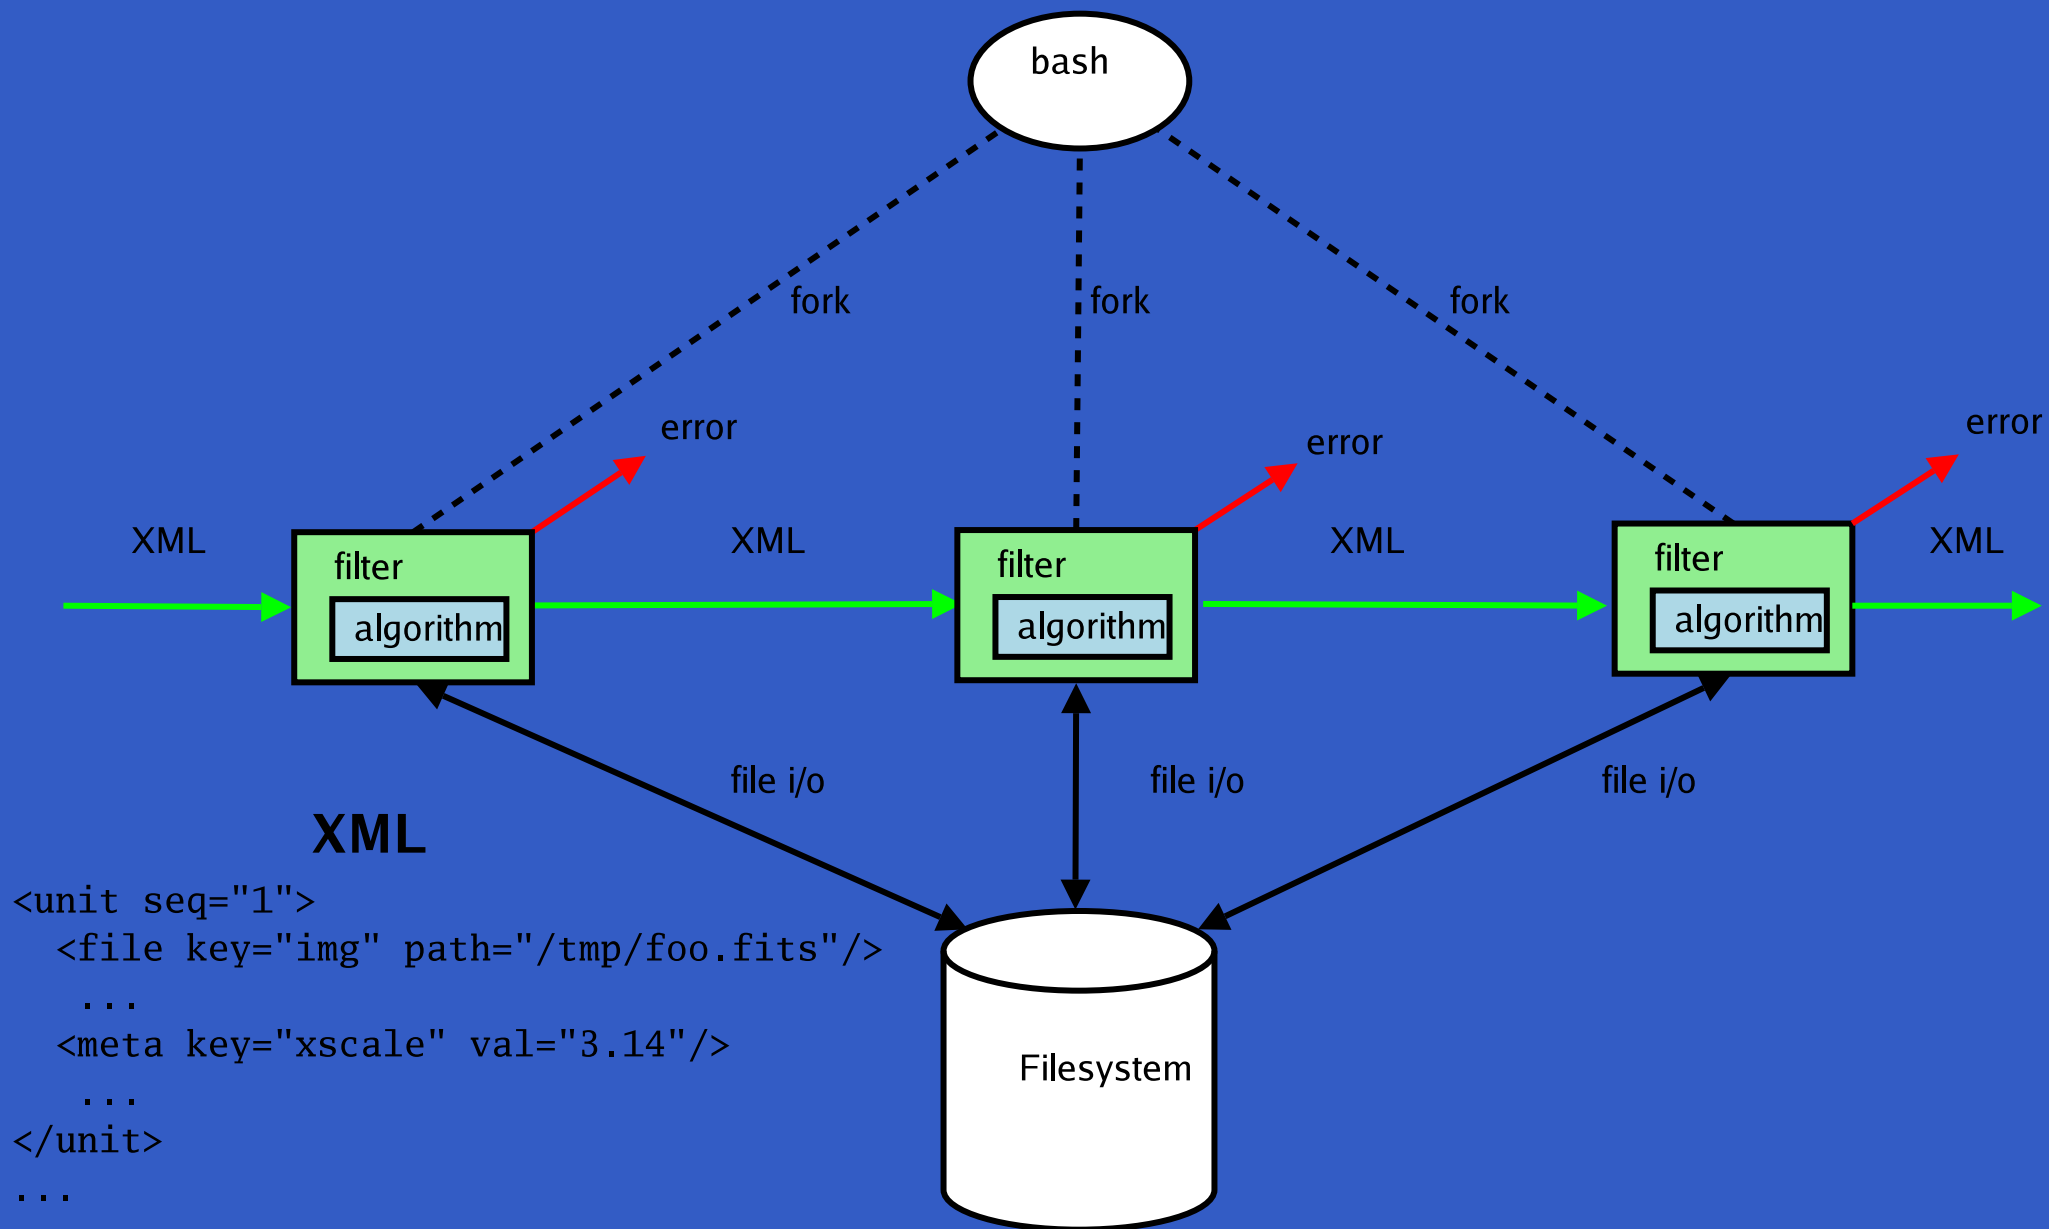


Overview

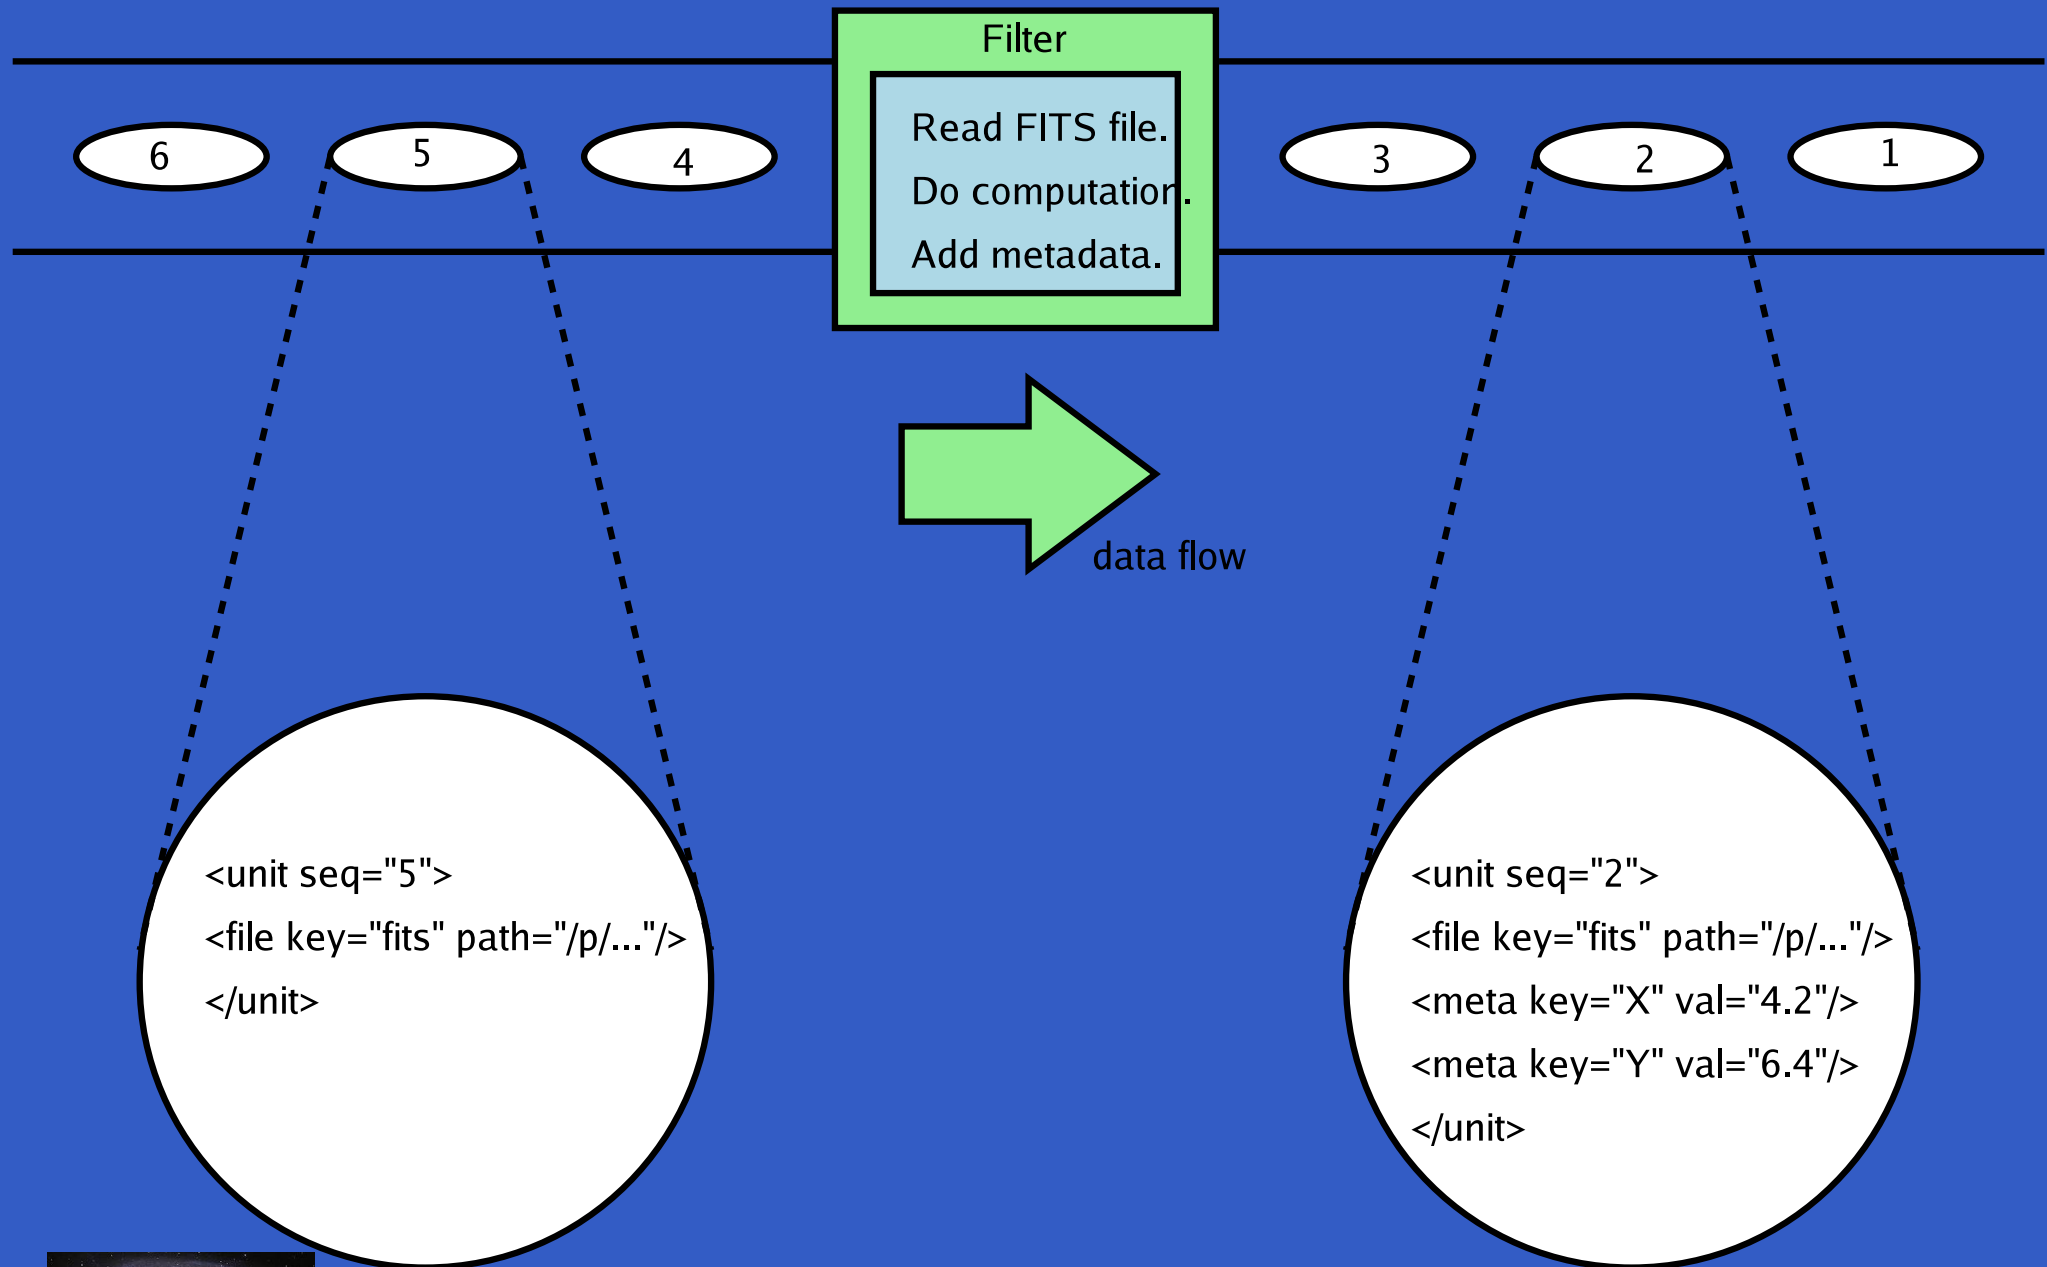
- Filters are linked to libisp.a, include isp.h.
- The shell and UNIX pipes tie filters together.
- Pipes carry XML stream of “work units”.
- Work units contain files (passed by reference).
- Work units contain metadata (passed by value).
- Filters produce, consume, or pass through data.



ISP Architecture



ISP Filters Process “Work Units”



ISP Pipelines are Shell Scripts

```
#!/bin/bash
```

```
filter [args] | filter [args] | filter [args] >results.xml
```

■ Good:


- ◆ Easy to build test harnesses, develop on laptop, etc.
- ◆ Existing UNIX filters can be used for debugging.
- ◆ Everybody already understands UNIX pipes.

■ Bad:

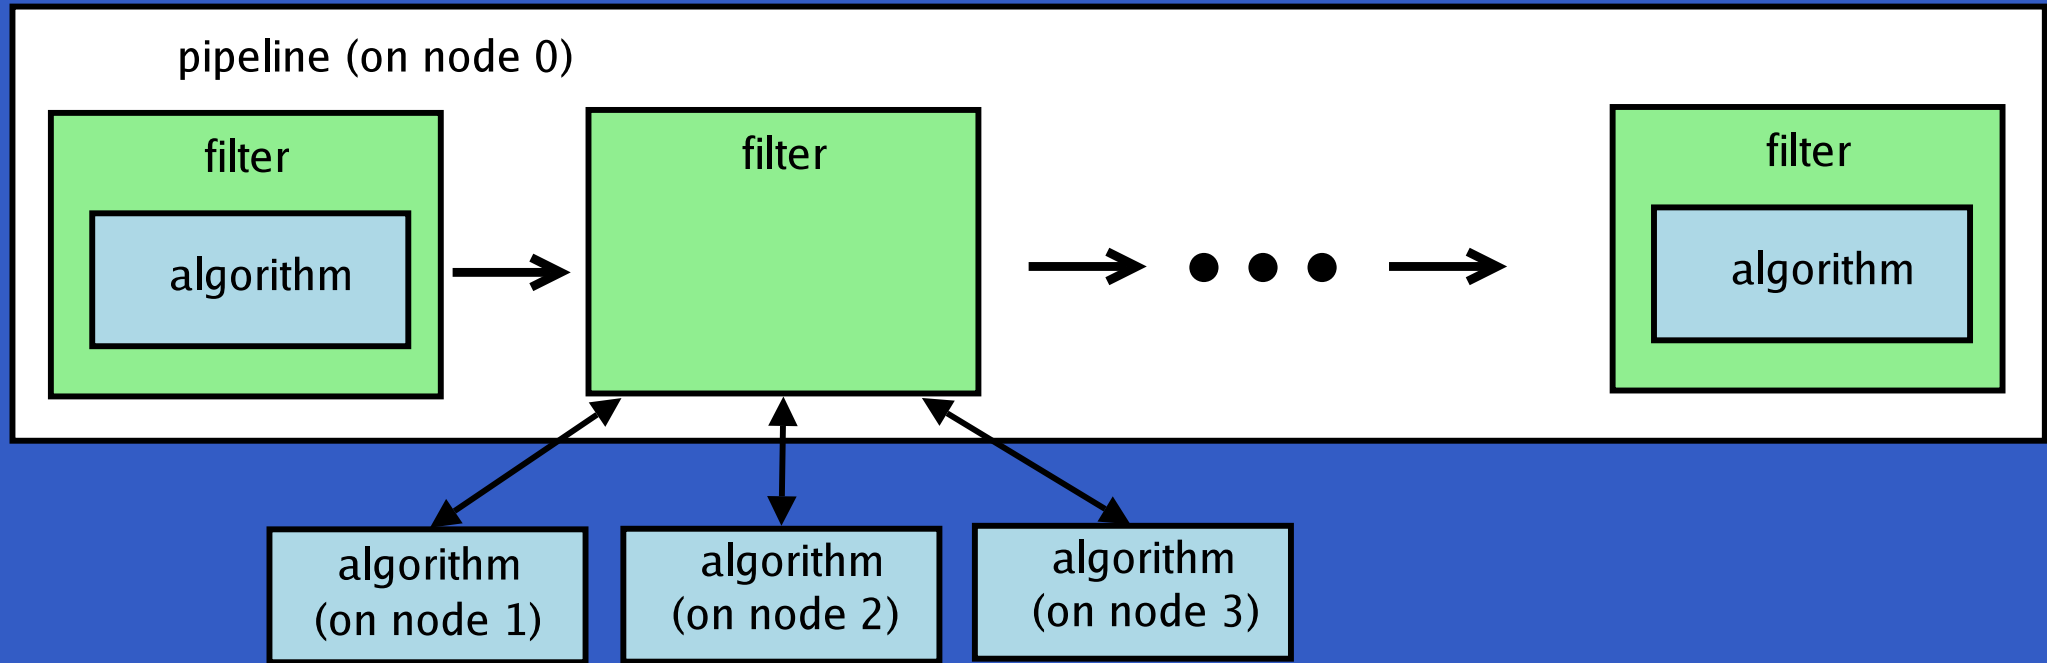
- ◆ There can be no cycles in the workflow.
- ◆ Forks and joins are awkward.



Or Filters Could Be Used in Other Systems



Work Units Can be Processed Concurrently



- Filters can offload work to a cluster resource manager.
- The details are not exposed to the filter programmer.
- Concurrency can be enabled/disabled per-filter.
- Future: Algorithm can use MPI.

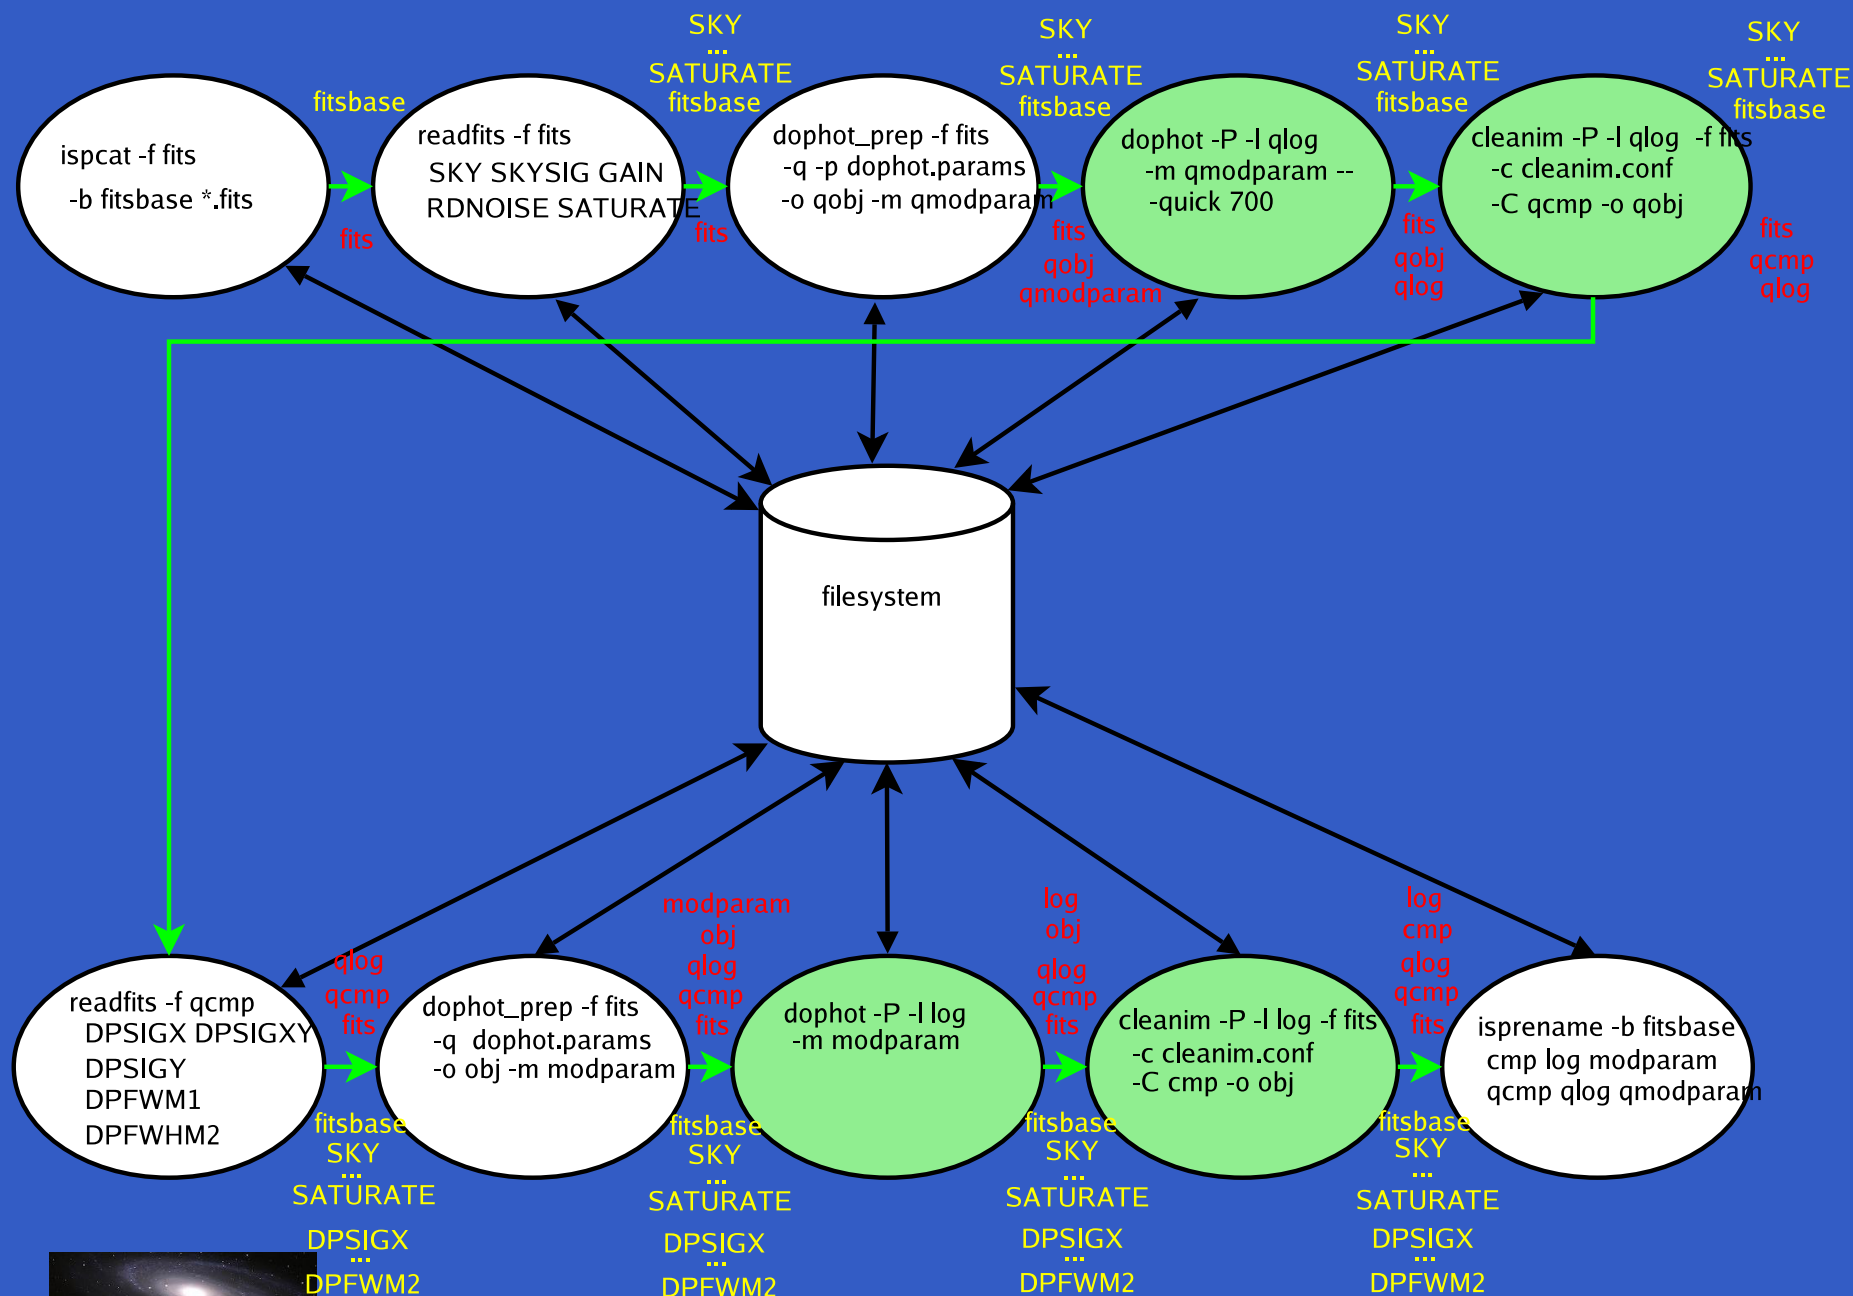


Example: SuperMACHO PhotPipe


- Implemented QUICKDOPHOT+CLEANIM; DOPHOT+CLEANIM stages.
- Did not implement FLATTEN stage.
- Some minor algorithmic stuff in perl was converted to additional “stages” in ISP.



PhotPipe with ISP



MCR PhotPipe Scaling Study



Version 0.9 Status

- C API bindings are provided.
- Sequential and SLURM data-parallel execution are supported.
- Next: convert more of SuperMACHO pipeline; make a public release.



References

For more information on ISP:

- ADASS 2004: *Middleware for Astronomical Data Analysis Pipelines* <http://www.llnl.gov/CASC/request/pubs.html>
- Presentation: *Industrial Strength Pipes: a pipeline toolkit for image processing*
<http://www.llnl.gov/CASC/request/presentations.html>

Citations:

- *MapReduce: Simplified Data Processing on Large Clusters*, Dean et al, <http://labs.google.com/papers/mapreduce.html>
- *GridDB: A Data-Centric Overlay for Scientific Grids*, Liu et al, <http://www.cs.berkeley.edu/~dtliu>



Disclaimer and Auspices

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

UCRL-PRES-209332

